2002-9

Digital Image Similarity for Geo-spatial Knowledge Management

James Carswell

Technological University Dublin, james.carswell@tudublin.ie

Recommended Citation

Carswell, J. (2002) Digital image similarity for geo-spatial knowledge management. *6th European Conference on Case-Based Reasoning (ECCBR2002),* Springer Verlag Lecture Notes in Artificial Intelligence. Aberdeen, Scotland. September, 4-7September.

2002-09-01

Digital image similarity for geo-spatial knowledge management

James D. Carswell
Dublin Institute of Technology, jcarswell@dit.ie

Digital Image Similarity for Geo-Spatial Knowledge Management

James D. Carswell[1], David C. Wilson[2], and Michela Bertolotto[2]

[1] Digital Media Centre, Dublin Institute of Technology, Dublin 2, Ireland
jcarswell@dit.ie
[2] Smart Media Institute, Department of Computer Science, University College Dublin, Belfield, Dublin 4, Ireland
{david.wilson,michela.bertolotto}@ucd.ie

Abstract. The amount and availability of high-quality geo-spatial image data, such as digital satellite and aerial photographs, is increasing dramatically. Task-based management of such visual information and associated knowledge is a central concern for organisations that rely on digital imagery. We are developing geo-spatial knowledge management techniques that employ case-based reasoning as the core methodology. In order to provide effective retrieval of task-based experiences that center around geo-spatial imagery, we need to forward novel similarity metrics for directly comparing the image components of experience cases. Based on work in geo-spatial image database retrieval, we are building an effective similarity metric for geo-spatial imagery that makes comparisons based on derived image features, their shapes, and the spatial relations between them. This paper gives an overview of the geo-spatial knowledge management context, describes our image similarity metric, and provides an initial evaluation of the work.

1 Introduction

Advances in sensor/scanner technology have resulted in the constantly increasing volume and availability of geo-spatial datasets, such as collections of digital satellite and aerial photographs. Moreover, the available imagery is becoming more complex, depicting characteristics of the earth surface and topography that are only visible in the near-infrared or microwave spectrum. The geosciences and spatial information engineering have been greatly affected by this information explosion. Geo-spatial information systems, in particular, have become crucial for addressing the problem of visual information overload by delivering on-point geographic image data combined with relevant associated information, and they play a key role in supporting the overarching task-based needs of organisations that rely on such information. Moreover, as geo-spatial information systems are used to address specific tasks, the expert interactions, analyses, and conclusions—based on relevant visual information—come to represent a substantial organizational knowledge asset.

For example, a company that uses geo-spatial data for architectural development projects may employ such a system to assist in selecting the optimal location for a new hanger at a major airport. From a task-based standpoint, the most relevant work product lies not merely in the applicable visual data, but in descriptions of why and how the information has been collected and to what ends it has been successfully (or unsuccessfully) employed. A clear advantage is provided by capturing and leveraging not only essential underlying information but also a measure of the human expertise involved in seeking out, distilling, and applying the information required for organisational tasks. This serves both to facilitate workflow by providing access to best-practice examples, as well as to grow a repository of task-based experience as a resource for support, training, and minimizing organisational knowledge-loss as a result of workforce fluctuations.

As part of an overall effort in intelligent geo-spatial information systems, we are developing case-based knowledge management support for libraries of geo-spatial imagery. The research draws on a substantial body of work in case-based knowledge management [39, 35, 18, 46, 6, 7, 32]. Our approach addresses task-based geo-spatial knowledge management by providing:

- digital image libraries for effective data organisation and efficient transmission to distributed clients
- sketch-based user interaction to provide a more natural mode of interaction in describing the context for retrieval
- a flexible task environment to support analysis and elucidation of relevant geo-spatial image information that can easily be integrated as part of existing workflow
- case-based tools to support intelligent capture and re-use of encapsulated task-based interactions and context

The challenges in the work are to integrate and tailor existing case-based methods to address specific needs for geo-spatial image information management, as well as to develop hybrid similarity measures that seamlessly integrate very different types of contextual knowledge afforded by query sketches, result images and metadata, image annotations, textual rationale, and other potential resource annotations (e.g., user voice/video recordings).

In order to provide effective retrieval of task-based experiences that center around geo-spatial imagery, we need to forward similarity metrics for directly comparing query sketches and image components of experience cases. Thus, in the first stage of the work, we are adapting and refining techniques developed for geo-spatial image database retrieval for use in the case-based components of the overall system. This paper describes our initial case-based similarity metric for geo-spatial imagery, which makes comparisons based on derived image features, their shapes, and the spatial relations between them. It is a straightforward derivation of work developed for image database retrieval [9, 2, 5]. Section 2 provides background on geo-spatial imagery, image retrieval, and integrations of CBR with imagery and GIS, and section 3 gives an overview of the task-based image retrieval context. Section 4 describes our approach to image indexing, while sections 5 and 6 respectively describe shape-based and relational components of image similarity. The paper goes on to describe the combined image similarity metric with a working example image query, and it concludes with brief discussion of future directions.

2 Background

This research draws on background in geo-spatial imagery, general approaches to image retrieval, and integrations of CBR with imagery and GIS.

2.1 Geo-spatial Imagery

Geo-spatial information represents the location, shape of, and relationships among geographic features and associated artifacts, including map and remotely sensed data. Two different formats are generally used to represent geo-spatial information: *raster* (digital images, with spatial position implicit in pixel ordering) and *vector* (layered coordinate representations with topographic and associated information, such as geographic maps and digital terrain models). In this research, we focus on managing the large quantities of geo-spatial information available in raster format, primarily digital aerial photos, satellite images, and raster cartography. Geo-spatial imagery is employed in a wide range of applications, such as intelligence operations, recreational and professional mapping, urban and industrial planning, and touristic systems. Typically, geo-spatial imagery will also include *metadata* information, such as: date and time of image acquisition; date and time of introduction to system; scale/resolution; location of the image, expressed in hierarchically arranged geographic entities (e.g., state, country, city); sensor information; and imagery type (e.g., black & white, colour, infrared).

2.2 Image Retrieval

Substantial research efforts within the computer vision community have been focused on retrieving specific images from a large database by querying the properties of these images [10, 25, 37, 41, 45]. Some notable prototypes for intelligent image retrieval have been developed, including [8, 17, 16, 19, 42, 47]. Most of these efforts address the problem in the context of general-use applications, where the images stored in the database display substantial differences in their low-level properties, such as: colour (histogram matching), texture (image coarseness and contrast matching), and composition (dividing an image into homogeneous colour/texture regions and analysing the relative positions of those regions).

An inherent characteristic with geo-spatial images, however, is that they are usually very similar in terms of general low-level properties. Thus in geo-spatial applications, image retrieval approaches based on low-level properties are not very effective. In geo-spatial applications, images are better distinguished by the shape and spatial configuration of the objects they contain. Consequently, a better approach to measuring similarity in geo-spatial datasets relies on the use of queries based on these higher-level properties.

2.3 Case-Based Reasoning

A number of research efforts have investigated case-based reasoning as applied to tasks involving imagery, such as medical diagnosis [36, 22, 44], face recognition [40], architectural support [11], protein crystallization [29], and remotely sensed data [48]. Previous research in case-based reasoning has examined image recognition [38] and segmentation [43]. For an overview of the issues involved in integrating imagery with case-based reasoning, see [20]. Case-based reasoning has also been applied in sketch-based retrieval of architectural data [12, 23], as well as for prediction in GIS applications [31, 30, 33, 27, 26].

Many of these case-based approaches rely on low-level image properties that are not appropriate for geo-spatial imagery. The spatial component in our domain also implies that there should not been any processing applied to the imagery (e.g., Fast Fourier Transforms) that transforms the raster image into the frequency domain before further operations begin. Closest in spirit to our work is [28], in which edge-image representations are used to index satellite imagery. While we plan to integrate some of the general techniques described in previous CBR research where applicable, we have chosen to base our image similarity metric on established work that defines measures tailored for geo-spatial imagery [9, 2, 5].

3 Task-Based Image Retrieval & Knowledge Management

A typical task-based query to our image repository is a straightforward request to a geo-spatial image database, and it could consist of specified metadata, semantic information, and a sketched configuration of image-objects [3, 4]. The metadata criterion would include such information on image scale or location, while the semantic criterion would match against previously entered annotations (if any) about the type (purpose, etc.) of objects that should be contained within images of interest. The sketch would include information on desired image-objects and their configuration. For example, if the user decided to retrieve all images with airplanes, airplane hangers, and runways that match to a particular configuration, the query would:

– Process the metadata to retrieve all images that match to the specified criteria (e.g., images from Dublin).
– From this subset of images, use any available semantic information (e.g., airplanes, terminal) to further constrain the result set.
– From this subset of images, select imagery indexed by object-features that best match the user sketch.
– Process the spatial relations of the sketch scene on the last image subset and return a prioritized list of imagery as the query result.

The task-based image retrieval tools under development are an effective means for locating geo-spatial image information, and they provide the core of the overall system. Alongside, we are developing tools for direct image manipulation, such as filters, transformations, highlighting, sketching, and post-it type annotations. These will allow the user to identify regions of interest that can be linked to clarifications, rationale, and other types of annotations (e.g., multimedia). The manipulations and annotations will not alter the underlying images or geo-spatial information, rather they will be layered to provide a task-specific view. This enables the capture and refinement of more general task-based ideas and rationale. A typical interaction with the system, then, can capture the sketch and geo-spatial query or queries posed by the user, the results that were found to be useful, as well as the user's annotations of the results. All of the contextual knowledge required to address the task goal can thus be captured as an experience case, enabling an increasingly powerful cycle of proactive support, with case-based suggestions based on task context.

As part of case-based retrieval of task experiences, we need to define a measure of similarity for directly comparing image components of experience cases, one which works in conjunction with the task-based image retrieval system. In doing so, we focus on the last two steps of the task-based retrieval.

4 Image Indexing

Image metadata and semantic information are used as part of the overall image indexing scheme, but from the standpoint of computing image-level similarity there are two main indexing dimensions: the edge-image representation and the feature library representation [9].

4.1 Image Pre-processing

Upon insertion into the image library, images are pre-processed automatically by first applying a high-pass edge enhancement filter and then applying a binary threshold to the resulting image, such that only black or white pixels remain. Spurious edges (of insignificant length) are also then removed by an additional filtering step. This process produces the edge-image representation (of a given raw image) on which shape similarity is computed. The edge-image (e.g., Figure 1b) thus contains only the boundary outlines of image-objects inherent to its corresponding raw image (e.g., Figure 1a). The original raw image is then stored along with its corresponding edge representation image.

4.2 Image Feature Library

The image feature library is a hierarchical arrangement of distinct feature outlines (i.e. image-object shapes) with links to image files where such features appear. It can be likened to an inverted term index in collaborative filtering. At the task-based level of retrieval, the feature library is used to reduce the search space of a query from the entirety of a large image library to a substantially reduced image set containing an abridged group of object-features.

From the standpoint of image-level similarity, the feature library defines a reduced vocabulary of canonical image-objects, subsets of which are used as

(a) (b)

Fig. 1. Raw and Edge Representation Images.

image indices. This index is used as a proxy for the actual image in computing image-level similarity. The individual feature outlines that comprise the library are smaller edge-images representing individual features of interest. They are derived through interaction with the system, either from user query sketches or explicitly identified image regions of interest. The feature library is itself an internal case-based component [34], but a full description is beyond the scope of this paper.

When a new image is inserted into the image library, all of the features in the feature library are matched (using the similarity metric defined in section 7) against the new image to see if and where they match. If a feature-to-image match is above a system threshold, a link from this feature to this image (and vice versa) is established along with its location (i.e. the coordinates of the feature centroid within the image), and the coordinates of the minimum bounding rectangle (MBR) of this feature within the image.

When a new query feature is added to the feature library, it is matched against the entire image library to establish indexing links. Because the indexing process can be computationally expensive, index maintenance is an off-line process.

5 Matching Image-Objects: S_{sh}

An image is composed of spatial objects and their relations with the similarity between scenes described as a function of object similarity plus relation similarity [21]. Our similarity metric is based on a measure of object shape similarity (S_{sh}) and three measures of relational similarity: topology (S_{top}), orientation (S_{or}), and distance (S_{dist}).

The method we employ to match image-object features is derived from least-squares "area-based" matching [1], which involves the extraction and matching of conjugate patches of pixels according to a correlation of summed squared pixel gray-level density differences [24]. In essence, the query patch slides across the library image (translating, rotating, and scaling), until a best-matching position is found.

In contrast to the gray-level matching in traditional least-squares methods, we reduce the comparison to essential shape information content by considering only those pixels that carry image-object information (i.e., image-object edges). This is facilitated by the binary edge-image representation and provides for good shape matching at a substantially reduced computational cost.

In our similarity metric, individual query image-objects are matched against a library image, and the results are combined in the overall image-level similarity measure. Given an image-object query (from a user sketch or the feature library), its number of rows and columns are noted along with the total number of pixels representing edges. The centroid of the image feature is calculated using the center of mass. The coordinates of the centroid pixel are then used as the origin for translation, rotation, and scaling during matching.

In order to account for local maxima, the library image is divided into a parameterized number (9 in practice) of disjoint regions, of equal or larger size to the query patch. The query patch is then matched within each of these regions in turn. In order to account for a feature being split across these regions, the query patch is also matched within each of the (16) overlap regions between the original divisions.

In matching a region, the query patch is divided into quadrants by its centroid. Each quadrant, then, measures the degree of match to its current location in the library image region by the extent of overlap in edge-pixels. Each edge-pixel in the quadrant contributes a vote either to stay (if it is already overlaps an image edge) or to move a certain distance in one of the cardinal directions (if it does not overlap). Move distance and direction are determined by the closest image edge. The individual pixel votes for direction and distance are summed, with higher weights given to shorter distances, to determine an overall shift for each quadrant.

From analysis of the edge pixel voting patterns, a decision is made to translate (quadrant votes in the same direction), scale (opposite quadrants vote in opposite directions), or rotate (quadrant votes follow a circular pattern) the query patch within the image region in order to acquire a better match. Typical distance values range from 0 to 10s of pixels in any direction with the maximum distance allowed being half the dimension of the query patch itself. Thus initial approximations for positioning the patch within the region are not required; if the patch is not in a suitable position within the image region, it will automatically move to the image content. This method also allows for occlusions of up to half of the query patch to be detected, as the patch can shift its origin (centroid pixel) right up to the border of the edge-image.

Once the patch is shifted into a new position, the process is repeated. Similar to the traditional least-squares approach, the solution is obtained after a set of iterations with parameterized boundary conditions on goodness of match (number of votes to stay exceed votes to move) and number of iterations (e.g., 20). The best-matched position for each library image region is recorded and the overall is used as the final matching position and percentage for the query patch in the library edge-image. When the query patch has settled on a match, its accuracy is determined by its matching percentage (i.e., by how many of its pixels continue to vote to stay put compared to the total number of pixels that constitute its edges).

6 Matching Image-Object Scenes

When query image-objects are matched to an image, their centroid coordinates within the image are recorded as well as the top left and bottom right coordinates of the query feature's minimum bounding rectangle, after scaling and rotation have taken place. Similarity for spatial relations on the image are determined through the use of the matched query features MBRs.

6.1 Matching Topology: S_{top}

Perhaps the most important of all spatial relations from a user's perspective is topology [14]. It is often more important, in composing a spatial query sketch, for a user to show that objects are positioned correctly relative to each other (disjoint, touching, overlapping, etc.) than to show their relative sizes or distances. It has been shown that topological relations can be derived automatically between pairs of simply connected regions (i.e., regions without holes, by determining their 4 intersection relations between their respective borders and interiors) [13]. More specifically, a 2x2 matrix is generated through the determination of whether the border b or interior i of region A intersects either the border or interior of region B. Taking into consideration the inconsistent relations, we are left with only 8 possible relations, shown in Figure 2. The 4 intersection method can be extended to describe a scene of n objects by building a nxn connectivity matrix whose elements consist of the relations between individual pairs of objects in the scene [15]. The result of such an operation gives a mathematical description of the topology of a scene that can be queried against for similarity. In defining our similarity measure, the degree of topological match is computed as a function of the number of steps between topological relation types (e.g., *disjoint*→*meets* is one step).

6.2 Matching Orientation: S_{or}

To overcome the lack of exterior orientation information in the query and image scenes, we use an intrinsic reference frame where the query/image-object orientation is in respect to *left-of* or *right-of* the features themselves. To do this, a

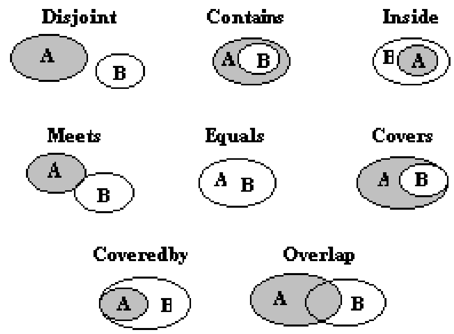

Fig. 2. Binary Topological Relations for Simply Connected Regions.

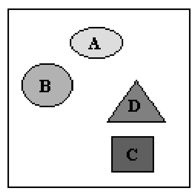

	A	B	C	D
AB	0	0	1	1
AC	0	-1	0	0
AD	0	-1	-1	0
BC	1	0	0	1
BD	1	0	-1	0
CD	1	1	0	0

Fig. 3. Example Image Scene and Corresponding Position Relation Matrix.

position relation matrix is built for each query/image scene. Since a scene comprised of two (or less) objects is trivial to discern (i.e., a scene could be rotated to suit any orientation), our approach assumes three (or more) scene objects. To build the position relation matrix for the scene depicted in Figure 3, an imaginary line connecting the centroid of Feature A to the centroid of Feature B is drawn. Feature A is arbitrarily considered as the *top* feature and Feature B the *bottom* feature. For every other feature (C through n) in the scene it is determined whether they lie *left-of* or *right-of* this line. Fixing the same features in both the query and image scenes to be either top or bottom renders any rotations in the scenes immaterial. The calculation of *left-of* or *right-of* is straightforward, given that we know (from the shape matching algorithm) the pixel coordinates of each feature's MBR in the image scene and each features actual boundary outline in the query scene. For example, if the MBR of Feature C in the image scene lies *left-of* this line, a value of -1 is placed in the position relation matrix. If Feature C's MBR lies *right-of* this line, a value of 1 is entered in the matrix and if Feature C's MBR lies somewhere along this extended line, a 0 is placed in the matrix.

After the remaining features of the scene are likewise added to the matrix, the imaginary line between Feature A and Feature B is deleted and redrawn between Feature A and Feature C with the process of determining *left-of* and *right-of* for

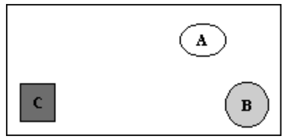

$$\begin{array}{c|ccc|} & AB & AC & BC \\ AB & 1 & 1.74 & 2.3 \\ AC & .57 & 1 & 1.32 \\ BC & .44 & .76 & 1 \end{array}$$

Fig. 4. Example Image Scene and Corresponding Distance Ratio Matrix.

the remaining features in the scene repeated. This entire procedure is repeated for every combination of features in the scene, i.e. there are $n(n-1)/2$ combinations of extended, imaginary lines that need to be tested against. For example, in the image scene of Figure 2, with 4 image-objects, there are 6 combinations of imaginary lines to be tested against, resulting in a 6x4 position relation matrix.

A position relation matrix is constructed for the query scene and image scene. A normalized correlation coefficient for the query/image scene is then calculated using Equation 1, to describe their respective similarities. This coefficient is scaled between 0 and 1 to give a total scene position matching percentage between the query and image-object configuration, independent of any arbitrary scene rotations.

$$NC = \frac{\sum((I-\overline{I})\cdot(Q-\overline{Q}))}{\sqrt{\sum(I-\overline{I})^2\cdot\sum(Q-\overline{Q})^2}} \tag{1}$$

I and Q are the position relation matrices for the image and query scenes respectively. Similar to the extrinsic reference frame approach, distance is also required in order to distinguish properly between these configurations of spatial entities.

6.3 Matching Relative Distance: S_{dist}

Where no scale information on the query/image scene is provided a-priori, it is necessary to analyze the relative distances between image-objects. A square matrix of rank $n(n-1)/2$ (where n is the number of objects in the scene) can be built for every query/image scene. Each combination of image-object connection is set to a distance of 1 unit, and the distances to the other objects in the scene determined relative to this unit distance. Assuming an example scene of 3 objects, a 3x3 distance ratio matrix would be built (Figure 4).

The ratio approach to calculating the relative distances between objects does not require absolute scale information and is possible of course because the pixel coordinates of the image-object centroids are returned from the shape matching algorithm. The similarity between various query and image scenes, or more specifically their respective distance ratio matrices, is then determined through analysis of their normalized correlation coefficients (Equation 1) scaled between 0 and 1.

7 A Similarity Metric for Image Scenes

Our image-level similarity metric combines similarity measures for individual object shapes, as well as the topological, orientation, and relative distance relations between them [2]. We define a similarity function S that assesses the similarity between an image query scene Q and an image scene I in the image library as follows:

$$S(Q,I) = S_{sh}(Q,I) \cdot w_{sh} + S_{top}(Q,I) \cdot w_{top} + \\ S_{or}(Q,I) \cdot w_{or} + S_{dist}(Q,I) \cdot w_{dist} \tag{2}$$

where:

- S_{sh} measures the degree of shape similarity between objects in Q and the corresponding objects in I. For example, assuming that $obj_1 \ldots obj_n$ are the objects in Q,

$$S_{sh}(Q,I) = \frac{\sum_{i=1}^{n} match\%(obj_i)}{n} \tag{3}$$

  where $match\%(obj_i)$ is the matching percentage between object obj_i in Q and the corresponding object in I. We might further constrain the match by imposing that for each $i = 1 \ldots n$ $match\%(obj_i) > \epsilon$, with a given threshold ϵ;
- S_{top} measures the degree of similarity between the set of topological relations characterizing objects in Q and the topological relations among corresponding objects in I;
- S_{or} measures the degree of similarity between the set of orientation relations characterizing objects in Q and the orientation relations among corresponding objects in I;
- S_{dist} measures the degree of similarity between the set of distance relations characterizing objects in Q and the distance relations among corresponding objects in I;
- w_{sh}, w_{top}, w_{or}, and w_{dist} are positive weights that establish the relative importance of the individual similarity criteria; their sum must equal 1.

7.1 A Working Example

In this section we show an example of how four different query scenes (Figure 5), sketched by the user, match to a given image I (Figure 6). The query scenes comprise differing configurations of the same three object shapes; i.e. an outline of an airplane (obj_1), airplane hanger (obj_2), and runway (obj_3). A summary of the results of the similarity metric calculations for this example can be found in Table 1. For each of the four query scenes, Q_1, Q_2, Q_3, Q_4, we calculate the value of S using Equation 2.

From these results it can be seen that Query Scene 2 is the best-matched configuration for the given image. This agrees with what a human observer would

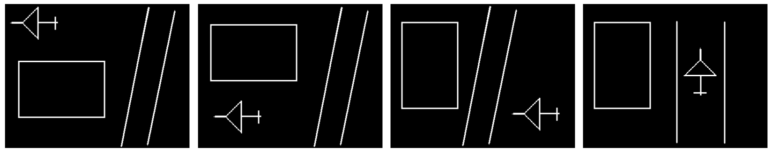

Q1 Q2 Q3 Q4

Fig. 5. Four Query Scenes.

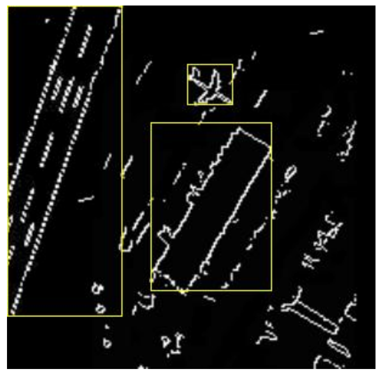

Fig. 6. Edge-Image with Superimposed MBRs.

choose as the best matched configuration since Query Scene 2 is plainly a rotation of the image, sketched at a significantly reduced scale. This demonstrates the ability of our approach to differentiate between arbitrary rotations and scaling of varying query/image scenes in addition to the capacity to distinguish between configurations and shapes of individual image-objects.

	Query Scene			
	1	2	3	4
S_{sh}	82.3	82.3	82.3	82.3
S_{top}	100	100	100	66.7
S_{or}	.5	100	79	81
S_{dist}	81.2	86.7	33.7	35
S	66	92.3	73.8	66.3

Table 1. Similarity Results for the Four Query Scenes

8 Conclusion

We have introduced a case-based reasoning approach to knowledge management in the context of task-based geo-spatial imagery retrieval. As a foundation for the building the overall case-based knowledge management component, we have derived an effective image-level similarity metric for directly comparing the image components of experience cases. The similarity metric operates in the raster/spatial domain and uses the shape of single image-object features together with their topological, orientation, and distance relations as matching primitives. We went on to present a working example as a practical illustration of how the similarity metric evaluates four query scenes for a given image.

While our first priority is to fully realize case-based support at the geo-spatial task-based level, our long-term goals include an extension of the image-matching technique for temporal analysis. By relaxing object/relation constraints and analyzing matching percentages, we expect to develop the temporal change detection in areas such as: object elimination, changes in object shape, and change in location. With the increasing use of geo-spatial information both in professional and recreational contexts, we expect that case-based approaches will prove invaluable in ameliorating problems of visual information overload.

References

1. F. Ackerman. Digital image correlation: Performance and potential application in photogrammetry. *Photogrammetric Record*, 11(64):429–439, 1984.
2. P. Agouris, M. Bertolotto, J. D. Carswell, C. Georgiadis, and A. Stefanidis. Scale and orientation-invariant scene similarity metrics for image queries. *International Journal of Geographical Information Science*, 16, 2002.
3. P. Agouris, J. Carswell, and A. Stefanidis. An environment for content-based image retrieval from large spatial databases. *ISPRS Journal of Photogrammetry & Remote Sensing*, 54(4):263–272, 1999.
4. P. Agouris, J. Carswell, and A. Stefanidis. Sketch-based image queries in topographic databases. *Journal of Visual Communication and Image Representation*, 10(2):113–129, 1999.
5. Peggy Agouris, Michela Bertolotto, James D. Carswell, and Anthony Stefanidis. A scene similarity metric for matching configurations of image objects. In *Proceedings of the XIX ISPRS Congress*, Amsterdam, 2000.
6. D. Aha, I. Becerra-Fernandez, F. Maurer, and H. Muñoz-Avila, editors. *Proceedings of the AAAI-99 Workshop on Exploring Synergies of Knowledge Management and Case-Based Reasoning*. AAAI Press, 1999.
7. I. Becerra-Fernandez and D. Aha. Case-based problem solving for knowledge management systems. In *Proceedings of the Twelfth Annual Florida Artificial Intelligence Research Symposium*, pages 219–223. AAAI, 1999.
8. C. Carson, S. Belongie, H. Greenspan, and J. Malik. Region-based image querying. In *Proceedings of the IEEE Workshop on Content-Based Access of Image and Video Libraries*, pages 42–49, San Juan, Puerto Rico, 1997.
9. J. Carswell. *Using Raster Sketches for Digital Image Retrieval*. PhD thesis, University of Maine, Orono, Maine, USA, 2000.

10. S.D. Cohen and L.J. Guibas. Shape-based indexing and retrieval; some first steps. In *Proceedings of the 1996 ARPA Image Understanding Workshop*, volume 2, pages 1209–1212, 1996.
11. Carl-Helmut Coulon. Image retrieval without recognition. In *First European Workshop on Case-Based Reasoning (EWCBR'93), Posters and Presentations*, volume 2 of *SEKI-Report 93-12-2*, pages 399–402, 1993.
12. Ellen Yi-Luen Do and Mark D. Gross. Reasoning about cases with diagrams. In *American Society of Civil Engineers (ASCE) 3rd Congress on Computing in Civil Engineering*, pages 314–320, 1996.
13. M. Egenhofer and R. Franzosa. Point-set topological spatial relations. *International Journal of Geographical Information Systems*, 5(2):161–174, 1991.
14. M. Egenhofer and D. Mark. Naive geography, spatial information theory - a theoretical basis for gis. In A. Frank and W. Kuhn, editors, *Lecture Notes in Computer Science*, pages 1–15, 1995.
15. M. Egenhofer and J. Sharma. Topological consistency. In *Proceedings of the 5th International Symposium on Spatial Data Handling*, pages 335–343, 1992.
16. D.A. Forsyth et al. Finding pictures of objects in large collections of images. In *Proceedings of the ECCV 96 Workshop on Object Representation*, 1996.
17. M. Flickner et al. Query by image and video content: The QBIC system. *IEEE Computer*, 28(9):23–32, 1995.
18. Kurt Fenstermacher and Carsten Tautz, editors. *Proceedings of the ICCBR-01 Workshop on Case-Based Reasoning Approaches for Process-Oriented Knowledge Management*. 2001.
19. C. Frankel, M. Swain, and W. Athitsos. Webseer: An image search engine for the world wide web. Technical Report TR-96-14, Department of Computer Science, University of Chicago, 1996.
20. J. Glasgow and I. Jurisica. Integration of case-based and image-based reasoning. In *Proceedings of the AAAI-98 Workshop on Case-Based Reasoning*, 1998.
21. R. Goyal and M. Egenhofer. Cardinal directions between extended spatial objects. *IEEE Transactions on Knowledge and Data Engineering*, 2000.
22. Morten Grimnes and Agnar Aamodt. A two layer case-based reasoning architecture for medical image understanding. In *Proceedings of EWCBR-96*, pages 164–178, 1996.
23. M. Gross, C. Zimring, and E. Do. Using diagrams to access a case library of design. In *Proceedings of the 1994 International Conference on Artificial Intelligence in Design*, pages 129–144. Kluwer, 1994.
24. A.W. Gruen and E.P. Baltsavias. Geometrically constrained multiphoto matching. *Photogrammetric Engineering and Remote Sensing*, 1987.
25. V.N. Gudivada and V.V. Raghavan. Design and evaluation of algorithms for image retrieval by spatial similarity. *ACM Transactions on Information Systems*, 13(2):115–144, 1995.
26. A. Holt and G. L. Benwell. Case-based reasoning and spatial analysis. *Journal of the Urban and Regional Information Systems Association*, 8(1):27–36, 1996.
27. A. Holt and G. L. Benwell. Applying case-based reasoning techniques in GIS. *The International Journal of Geographical Information Science*, 13(1):9–25, 1999.
28. B. Jose, B. P. Singh, S. Venkataraman, and R. Krishnan. Vector based image matching for indexing in case based reasoning systems. In *Proceedings of GWCBR-96*, 1996.
29. I. Jurisica, P. Rogers, J. Glasgow, S. Fortier, J. Luft, D. Bianca, and G. DeTitta. Image-feature extraction for protein crystallization: Integrating image analysis and

case-based reasoning. In *Proceedings of the Thirteenth Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-2001)*, pages 73–80, 2001.
30. Daniel Kaster, Heloisa Vieira Rocha, and Claudia Bauzer Medeiros. Case-based reasoning applied to environmental modeling with GIS. In *Proceedings of GIScience 2000*, 2000.
31. A. Khattak and H. Renski. PLAN<>HOV: A case-based reasoning planning tool for high-occupancy-vehicle lane analysis in a GIS environment. *Transportation Research Record*, 1682:18–27, 1999.
32. H. Kitano and H. Shimazu. The experience sharing architecture: A case study in corporate-wide case-based software quality control. In D. Leake, editor, *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, pages 235–268. AAAI Press, Menlo Park, CA, 1996.
33. Yu Ting Hung Ko-Wan Tsou, Yao-Lin Chang. A case-based urban planning support system using an integrated combination of geographical information systems and remote sensing. In *Proceedings of the 21st Asian Conference on Remote Sensing*, 2000.
34. D. Leake. Case-based CBR: Capturing and reusing reasoning about case adaptation. *International Journal of Expert Systems*, 1997.
35. David B. Leake and David C. Wilson. A case-based framework for interactive capture and reuse of design knowledge. *Applied Intelligence*, 14(1), 2001.
36. R. T. Macura and K. J. Macura. MacRad: Case-based retrieval system for radiology image resource. In *Proceedings of ICCBR-95*, 1995.
37. R. Mehrotra and J. Gray. Similar-shape retrieval in shape data management. *IEEE Computer*, 28(9):57–62, 1995.
38. Alessandro Micarelli, Alessandro Neri, and Giuseppe Sansonetti. A case-based approach to image recognition. In *Proceedings of EWCBR-2000*, pages 443–454, 2000.
39. Mirjam Minor and Steffen Staab, editors. *Proceedings of the German Workshop on Experience Management*. 2002.
40. A. D. Narasimhalu. CAFIIR: An image based CBR/IR application. In *Proceedings of the 1993 AAAI Spring Symposium on Case-Based Reasoning and Information Retrieval*, pages 70–77, 1993.
41. V.E. Ogle. Chabot: Retrieval from a relational database of images. *IEEE Computer*, pages 23–32, September 1995.
42. A. Pentland, R. W. Picard, and S. Scarloff. Photobook: Content-based manipulation of image databases. *International Journal of Computer Vision*, 18(3):233–254, 1996.
43. P. Perner. An architeture for a cbr image segmentation system. *Engineering Applications of Artificial Intelligence*, 12(6):749–759, 1999.
44. P. Perner. CBR-based ultra sonic image interpretation. In *Proceedings of EWCBR-2000*, pages 479–481, 2000.
45. R.W. Pickard and T.P. Minka. Vision texture for annotation. *Multimedia Systems*, 3(1):3–14, 1995.
46. *Proceedings of GWCBR-01: Knowledge Management by Case-Based Reasoning: Experience Management as Reuse of Knowledge*. 2001.
47. S. Sclaroff, L. Taycher, and M. La Cascia. Imagerover: A content-based image browser for the world wide web. In *Proceedings of the IEEE Workshop on Content-Based Access of Image and Video Libraries*, pages 2–9, 1997.
48. S. Venkataraman, R. Krishnan, and K. Rao. A rule-rule-case based system for image analysis. In *First European Workshop on Case-Based Reasoning (EWCBR'93), Posters and Presentations*, volume 2 of *SEKI-Report 93-12-2*, pages 410–414, 1993.